Microsoft Ireland Operations Ltd
One Microsoft Place
South County Business Park
Leopardstown
DUBLIN D18 P521

7 November 2019

Department of Business, Enterprise and Innovation
23 Kildare Street
Dublin D02 TD30
Email: aistrategy@dbei.gov.ie

Re: Submission to the Department of Business, Enterprise and Innovation with respect to the Public Consultation on the Development of a National Strategy on Artificial Intelligence

To Whom It May Concern:

Microsoft respectfully submits the following comments in response to the Department of Business, Enterprise and Innovation's Public Consultation on the Development of a National Strategy on Artificial Intelligence.

In summary, we have contributed to and are fully supportive of the submission by Ibec for this consultation exercise. In addition, we have provided some further perspective in this response. We appreciate the opportunity to provide our input to this consultation and welcome follow-up discussions with the Department on our response below.

Respectfully submitted,

On behalf of Microsoft Ireland

by

[redacted]

/ctd.

Microsoft's submission to the Department of Business, Enterprise and Innovation with respect to the Public Consultation on the Development of a National Strategy on Artificial Intelligence

Executive Summary

Overview

Recognising the remarkable cluster of technology companies located here, there is a unique opportunity for Ireland to emerge as Europe's digital leader. This is particularly true at this time when European policymakers are seeking to foster greater growth within the EU's digital single market, and when technology policy is at the centre of the European agenda.

The traditional and very successful Irish foreign direct investment message of 'tax, talent, and track record' needs to be reinforced with additional attributes, such as innovation, digital transformation, and technology policy leadership.

Ireland's leadership on an issue like AI adoption, its use and regulation, should be seen as a considerable opportunity for Ireland to establish its digital leadership credentials on the European stage.

To that end, Microsoft believes that Ireland should lead the way in calling for and supporting the regulation of facial recognition technology and commit itself to working closely with the new Commission and incoming President von der Leyen on her pledge to introduce regulation in this area.

Microsoft AI Principles

At Microsoft, we believe that six principles should provide the foundation for the development and deployment of AI-powered solutions that will put humans at the center:

- Fairness: to ensure fairness, we must understand how bias can affect AI systems;
- Safety and Reliability: people should play a critical role in making decisions about how and when AI systems are deployed;
- Privacy and Security: AI systems must comply with privacy laws that regulate data collection, use, and storage;
- Inclusiveness: AI solutions must address a broad range of human needs and experiences through inclusive design;
- Transparency: people must understand how decisions are made and then can more easily identify potential bias, errors, and unintended outcomes; and
- Accountability: as in areas like healthcare and privacy, people who design and deploy AI systems must be accountable for how their systems operate.

Standards in AI

It is important to note that AI is not just a discrete information technology (IT), it is also co-dependant on many other IT frameworks such as cybersecurity, privacy, data protection and IT governance that inform IT use. Perhaps the most important part of that structured co-dependence is standards, created via the international standards system.

Compliance with a standard provides greater assurance to the consumer than a principle can. The international standards system can help reconcile the needs for compliance and innovation in a way that, we believe, can lead to improved understanding of the challenges and how to address them.

Facial Recognition Challenges

Facial Recognition Regulation

Microsoft has adopted six principles for facial recognition technology, and we would recommend that the Government considers these as it develops its policy in this area. These include: fairness, transparency, accountability, non-discrimination, notice and consent, and lawful surveillance.

Bias and Discrimination

Certain uses of facial recognition technology increase the risk of decisions, outcomes, and experiences that are biased and can even violate discrimination laws.

We believe new laws can address this need with a two-pronged approach:

1. Requiring transparency
   Legislation should require tech companies that offer facial recognition services to provide documentation that clearly explains the capabilities and limitations of the technology in terms that people can understand.

2. Enabling third-party testing and comparisons
   New laws should also require that providers of commercial facial recognition services enable third parties engaged in independent testing to conduct and publish reasonable tests of their facial recognition services for accuracy and unfair bias.

Additional Measure – Future Laws

The problems of bias and discrimination can be exacerbated when organisations deploy facial recognition beyond the limits of current technology or in a manner that is different from what was intended when the technology was designed.

We believe that a new law can help address these concerns without imposing onerous obligations on businesses and other users. This too involves a two-pronged legal approach:

1. Ensuring meaningful human review
   In certain high-stakes scenarios, it's critical for qualified people to review facial recognition results and make key decisions rather than simply turn them over to computers. New legislation should therefore require entities that deploy facial recognition to undertake meaningful human review of facial recognition results prior to making final decisions where decisions may create a risk of bodily or emotional harm to a consumer, where there may be implications on human or fundamental rights, or where a consumer's personal freedom or privacy may be impinged.

2. Avoiding use for unlawful discrimination
   It's important for entities that deploy facial recognition services to understand that they are not absolved of their obligation to comply with laws prohibiting discrimination against individual consumers or groups of consumers. This provides further reason to ensure that humans undertake meaningful review, given their ongoing accountability under the law for decisions based on the use of facial recognition.

Protecting People's Privacy
People deserve to know when facial recognition technology is being used, so they can ask questions and exercise choice in the matter if they wish.

This type of transparency is vital for building public knowledge and confidence in this technology and new legislation can provide for this in a straightforward approach:

1. Ensuring Notice
   The law should require entities that use facial recognition to identify consumers, place conspicuous notice that clearly conveys that these services are being used.
2. Clarifying Consent
   The law should specify that consumers consent to the use of facial recognition services when they enter premises or proceed to use online services that have this type of clear notice.

Protecting Democratic Freedoms and Human Rights
There are many governmental uses of facial recognition technology that will protect public safety and promote better services for the public without raising these types of concerns. But there is one potential use for facial recognition technology that could put our fundamental freedoms at risk and it is important to address this risk.

When combined with ubiquitous cameras and massive computing power and storage in the cloud, a government could use facial recognition technology to enable continuous surveillance of specific individuals. This use of facial recognition could unleash mass surveillance on an unprecedented scale.

An indispensable democratic principle has always been the tenet that no government is above the law. Today this requires that we ensure governmental use of facial recognition technology remain subject to the rule of law.

Microsoft believes that to protect against the use of facial recognition to encroach on democratic freedoms, legislation should only permit law enforcement agencies to use facial recognition to engage in ongoing surveillance of specified individuals in public spaces when:

- a court order has been obtained to permit the use of facial recognition services for this monitoring; or
- where there is an emergency involving imminent danger or risk of death or serious physical injury to a person.

1. The Need for Responsible Artificial Intelligence

Advances in Artificial Intelligence (AI) are giving rise to computing systems that can see, hear, learn and reason, creating new opportunities to improve education and healthcare, address poverty, and achieve a more sustainable future.

But these rapid technology changes also raise complex questions about the impact they will have on other aspects of society: jobs, privacy, safety, inclusiveness, and fairness. When AI augments human decision-making, how can we ensure that it treats everyone fairly, and is safe and reliable? How do we respect privacy? How can we ensure people remain accountable for systems that are becoming more intelligent and powerful?

To realize the full benefits of AI, we'll need to work together to find answers to these questions and create systems that people trust. Ultimately, for AI to be trustworthy, we believe that it must be "human-centered" –designed in a way that augments human ingenuity and capabilities –and that its development and deployment must be guided by ethical principles that are deeply rooted in timeless values. At Microsoft, we believe that six principles should provide the foundation for the development and deployment of AI-powered solutions that will put humans at the center:

1.1. FAIRNESS

When AI systems make decisions about medical treatment or employment, for example, they should make the same recommendations for everyone with similar symptoms or qualifications. To ensure fairness, we must understand how bias can affect AI systems.

1.2. SAFETY AND RELIABILITY

AI systems must be designed to operate within clear parameters and undergo rigorous testing to ensure that they respond safely to unanticipated situations and do not evolve in ways that are inconsistent with original expectations. People should play a critical role in making decisions about how and when AI systems are deployed.

1.3. PRIVACY AND SECURITY

Like other cloud technologies, AI systems must comply with privacy laws that regulate data collection, use and storage, and ensure that personal information is used in accordance with privacy standards and protected from theft.

1.4. INCLUSIVENESS

AI solutions must address a broad range of human needs and experiences through inclusive design practices that anticipate potential barriers in products or environments that can unintentionally exclude people.

1.5. TRANSPARENCY

As AI increasingly impacts people's lives, we must provide contextual information about how AI systems operate so that people understand how decisions are made and can more easily identify potential bias, errors and unintended outcomes.

1.6. ACCOUNTABILITY

People who design and deploy AI systems must be accountable for how their systems operate. Accountability norms for AI should draw on the experience and practices of other areas, such as healthcare and privacy, and be observed both during system design and in an ongoing manner as systems operate in the world.

2. Collaboration on AI

A continuing collaboration between government, business, civil society and academic researchers will be essential to shape the development and deployment of human-centered AI to be trustworthy. Ongoing dialogues among these communities will help to identify and prioritize issues of societal importance, enable further research and development of solutions and sharing of best practices as new issues emerge, and, where appropriate, shape policy that can more readily adapt to these rapidly evolving technologies.

Some of the most urgent questions surrounding AI relate to jobs and work. Today, there are clear indications that how businesses organise work, how people find employment, and the skills people need to prepare for the workforce are changing dramatically. These changes are likely to accelerate in the decade ahead.

AI and cloud computing are driving much of this change. This is evident in the rise of the on-demand –or "gig" –economy and the automation of many aspects of today's jobs. As AI continues to transform the nature of work, we'll need to think in new ways about education, skills and training to ensure that people are prepared for the jobs of the future and businesses have access to the talent they need to succeed. And as traditional models of employment transform, we'll need to modernize legal frameworks to recognize new ways of working, provide adequate worker protections and maintain social safety nets.

To enable people to thrive in today's economy, and prepare for tomorrow's, we believe it's critical to focus on the following areas:

2.1. Preparing Today's Students

Every young person should have the opportunity to study computer science. The skills they gain will open the door to higher-paying jobs in faster-growing fields. This means equitable access to rigorous and engaging computer science courses and a focus on uniquely human skills must be top priorities.

2.2. Supporting Today's Workers

We must also help today's workers gain skills that are relevant in the changing workplace. Distance and online learning and investments in on-the-job training programs will be essential. And we'll also need to improve how we identify the skills that businesses need.

2.3. Creating a Skills-based Marketplace

To help companies find qualified employees and enable workers to find jobs, we'll need to move from a degree-based system to one that uses credentials that are widely recognized and valued by employers.

2.4. Providing Legal Certainty for Employers and Workers

The rise of the on-demand economy raises important questions that are not clearly addressed by existing laws about how we classify workers. To enable innovation and to protect workers, legal certainty must be created so that workers and businesses understand their rights and obligations.

2.5. Developing Industry Standards to Protect Workers

Business leaders have an opportunity to play a significant role in reshaping employment policy in the emerging economy by setting their own standards for on-demand engagements that include fair pay and treatment for on-demand workers.

2.6. Ensuring Benefits Move with Workers and Modernising Social Safety Nets

As the nature of work evolves with technology innovation, the traditional system of employer-provided benefits and government-supported social safety needs to be reformed to provide adequate coverage for workers and a sustainable contribution structure for businesses.

3. Regulation on Facial Recognition

The President of the European Commission has indicated that new legislation on AI will be introduced. A number of EU countries, including Germany, have already adopted national strategies in this area and are now pushing for greater cross-border cooperation. Chancellor Angela Merkel is a vocal proponent of a legally binding body of European rules[1] to ensure that AI 'serves humanity' – something which Microsoft believes is urgently needed.

AI is already embedded in tools we use every day, from the navigation systems in our cars that show us how to avoid traffic jams, to the streaming services that suggest what movie to watch next. All of us make daily choices based on offers pre-curated by AI. The fact that these services make our lives so much easier means that we rarely question them. As beneficial as AI can be at both an individual and a societal level – from helping with early cancer detection and tackling climate change, to assisting law enforcement with the prosecution of serious crimes such as child pornography – it also raises serious questions about the protection of fundamental rights. This is particularly true when it comes to facial recognition technology.

Around the world, activists have been warning about the misuse of facial recognition technology for mass surveillance and censorship. In the UK, the increased use of facial recognition in shopping centres, museums and conference centres has been deemed an 'epidemic' by privacy campaigners. In Germany, a facial recognition pilot deployed in Berlin's Südkreuz railway station[2] has been fiercely criticised by lawyers and data protection groups.

The debate around facial recognition demonstrates the need for regulators to take the issue firmly in hand. This is not just about what the technology can do, but what it should do, and how it should do it. High-level principles[3] can provide a roadmap for an ethical approach, but they aren't enough. When it comes to facial recognition, we need binding regulation.

3.1. Facial Recognition Problems

First, especially in its current state of development, certain uses of facial recognition technology increase the risk of decisions and, more generally, outcomes that are

[1] https://www.bundesregierung.de/breg-en/news/merkel-zu-kuenstlicher-intelligenz-1639836, last accessed 30 October 2019.

[2] https://www.politico.eu/article/berlin-big-brother-state-surveillance-facial-recognition-technology/, last accessed 30 October 2019.

[3] https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai, last accessed 30 October 2019.

biased and, in some cases, in violation of laws prohibiting discrimination.

Second, the widespread use of this technology can lead to new intrusions into people's privacy.

And third, the use of facial recognition technology by a government for mass surveillance can encroach on democratic freedoms.

We believe all three of these problems should be addressed through legislation, as described below.

3.2. Addressing Bias and Discrimination

Certain uses of facial recognition technology increase the risk of decisions, outcomes, and experiences that are biased and even in violation of discrimination laws. Recent research has demonstrated, for example, that some facial recognition technologies have encountered higher error rates when seeking to determine the gender of women and people of colour. This makes it especially important that technology companies continue the work needed to identify and reduce these errors and improve the accuracy and quality of facial recognition tools and services. It's equally critical that we work with customers closely to ensure that facial recognition services are deployed properly in ways that will reduce these risks. Over time, we believe that well-functioning market forces can encourage the technology innovation that is needed.

But we also believe that new laws are needed in this area, and for two distinct reasons. First, market forces will work well only if potential customers are well-informed and able to test facial recognition technology for accuracy and risks of unfair bias, including biases that arise in the context of specific applications and environments. Technology companies currently vary in their willingness to make their technology available for this purpose. As a result, some academic tests of these services have omitted some of the market leaders. And when important advocacy organizations have tried to perform tests, they've almost immediately been met by rejections and criticism by some providers who claim that the testing is deficient. As a society, we need legislation that will put impartial testing groups like Consumer Reports and their counterparts in a position where they can test facial recognition services for accuracy and unfair bias in a transparent and even-handed manner.

We believe new laws can address this need with a two-pronged approach:

3.2.1. Requiring transparency

Legislation should require tech companies that offer facial recognition services to provide documentation that explains the capabilities and limitations of the technology in terms that customers and consumers can understand.

3.2.2. Enabling third-party testing and comparisons

New laws should also require that providers of commercial facial recognition services enable third parties engaged in independent testing to conduct and publish reasonable tests of their facial recognition services for accuracy and unfair bias. A sensible approach is to require tech companies that make their facial recognition services accessible using the internet also make available an application programming interface or other technical capability suitable for this purpose.

There's a second reason we believe new legislation is needed in this area now, and it points to additional measures that new laws should address. While we're hopeful that market forces may eventually solve issues relating to bias and discrimination, we've witnessed an increasing risk of facial recognition services being used in ways that may adversely affect consumers and citizens today.

It's obviously of little solace to think about the eventual improvements in this technology if it misidentifies you and is used in a way that deprives you of the ability to access government services, obtain admission to an event or purchase commercial products. These problems can be exacerbated when organizations deploy facial recognition beyond the limits of the current technology or in a manner that is different from what was intended when they were designed. We believe that new laws can help address these concerns without imposing onerous obligations on businesses and other users. This too involves a two-pronged legal approach:

3.2.3. Ensuring meaningful human review

While human beings of course are not immune to errors or biases, we believe that in certain high-stakes scenarios, it's critical for qualified people to review facial recognition results and make key decisions rather than simply turn them over to computers. New legislation should therefore require that entities that deploy facial recognition undertake meaningful human review of facial recognition results prior to making final decisions for what the law deems to be "consequential use cases" that affect consumers. This includes where decisions may create a risk of bodily or emotional harm to a consumer, where there may be implications on human or fundamental rights, or where a consumer's personal freedom or privacy may be impinged.

3.2.4. Avoiding use for unlawful discrimination

Finally, it's important for the entities that deploy facial recognition services to recognize that they are not absolved of their obligation to comply with laws prohibiting discrimination against individual consumers or groups of consumers. This provides additional reason to ensure that humans undertake meaningful review, given their ongoing and ultimate accountability under the law for decisions that are based on the use of facial recognition.

3.3. Protecting People's Privacy

The widespread use of facial recognition technology can lead to new intrusions into people's privacy. For example, every public establishment could install cameras connected to the cloud with real-time facial recognition services.

Interestingly, the privacy movement in the United States was born from improvements in camera technology. In 1890, future Supreme Court Justice Louis Brandeis took the first step in advocating for privacy protection when he co-authored an article with colleague Samuel Warren in the Harvard Law Review advocating "the right to be let alone." The two argued that the development of "instantaneous photographs" and their circulation by newspapers for commercial gain had created the need to protect people with a new "right to privacy."

Technology today gives a new meaning to "instantaneous photographs" that Brandeis and Warren probably never imagined. From the moment one steps into a shopping mall, it's possible not only to be photographed but to be recognized by a computer wherever one goes. Beyond information collected by a single camera in a single session, longer-term histories can be pieced together over time from multiple cameras at different locations. A mall owner could choose to share this information with every store. Stores could know immediately when you visited them last and what you looked at or purchased, and by sharing this data with other stores, they could predict what you're looking to buy on your current visit.

Our point is not that the law should deprive commercial establishments of this new technology. On the contrary, we are among the companies working to help stores responsibly use this and other digital technology to improve shopping and other consumer experiences. We believe that a great many shoppers will welcome and benefit from improvements in customer service that will result.

But people deserve to know when this type of technology is being used, so they can ask questions and exercise some choice in the matter if they wish. Indeed, we believe this type of transparency is vital for building public knowledge and confidence in this technology.

New legislation can provide for this in a straightforward approach:

3.3.1. Ensuring Notice

The law should require that entities that use facial recognition to identify consumers place conspicuous notice that clearly conveys that these services are being used.

3.3.2. Clarifying Consent

The law should specify that consumers consent to the use of facial recognition services when they enter premises or proceed to use online services that have this type of clear notice.

3.4. Protecting Democratic Freedoms and Human Rights

The use of facial recognition technology by a government for ongoing surveillance of specified individuals encroaches on democratic freedoms and human rights. Democracy has always depended on the ability of people to assemble, to meet and talk with each other and even to discuss their views both in private and in public. This in turn relies on the ability of people to move freely and without constant government surveillance. Today this requires that we ensure that governmental use of facial recognition technology remain subject to the rule of law.

There are many governmental uses of facial recognition technology that will protect public safety and promote better services for the public without raising these types of concerns. There is an increasing number of such services in place already, and we should encourage them subject to the other protections described here.

New legislation can put us on this path:

3.4.1. Limiting Ongoing Government Surveillance of Specified Individuals

To protect against the use of facial recognition to encroach on democratic freedoms, legislation should permit law enforcement agencies to use facial recognition to engage in ongoing surveillance of specified individuals in public spaces only when: a court order has been obtained to permit the use of facial recognition services for this monitoring; or where there is an emergency involving imminent danger or risk of death or serious physical injury to a person.

It will be important for legislators to consider the standards for obtaining court orders in this area. For the most part, we believe this should be based on traditional rules like probable cause for search warrants. But there may be other narrow circumstances that also should be permissible, such as appropriate uses to help locate missing persons.

3.5. Looking Beyond Law and Regulation

While we believe that new laws and regulations are indispensable, we also recognise that they are not a substitute for the responsibility that needs to be exercised by technology companies. Microsoft has adopted six principles for facial recognition technology, based on the Principles for AI that we described above and we would recommend that the Government considers these as it develops its policy in this area.

These principles are:

3.5.1. Fairness

We will work to develop and deploy facial recognition technology in a manner that strives to treat all people fairly.

3.5.2. Transparency

We will document and clearly communicate the capabilities and limitations of facial recognition technology.

3.5.3. Accountability

We will encourage and help our customers to deploy facial recognition technology in a manner that ensures an appropriate level of human control for uses that may affect people in consequential ways.

3.5.4. Non-discrimination

We will prohibit in our terms of service the use of facial recognition technology to engage in unlawful discrimination.

3.5.5. Notice and consent

We will encourage private sector customers to provide notice and secure consent for the deployment of facial recognition technology.

3.5.6. Lawful surveillance

We will advocate for safeguards for people’s democratic freedoms in law enforcement surveillance scenarios and will not deploy facial recognition technology in scenarios that we believe will put these freedoms at risk.

4. Importance of Standards in AI

Microsoft is deeply committed to working with stakeholders to build a strong trust fabric underpinning AI which can give technical effect to the Microsoft Principles for AI.

4.1. AI and other IT Frameworks

It's important to note that AI is not just a discrete Information technology, it also co -depends on a mosaic of other extant IT frameworks such as cybersecurity, privacy, data protection and IT governance, and on other considerations that inform IT use, for example:

4.1.1. interoperability between different systems; and

4.1.2. applicability across different industries and sectors, for example, healthcare and banking.

Perhaps the most important part of that structured co-dependence is standards, created via the International standards system[4]. More generally, the standards produced via the international standards system can often provide important resources and references for both legislators and regulators as they adapt existing national IT policies to take account of AI scenarios.

Microsoft has been deeply engaged in standards development work across these domains of cybersecurity, data protection, privacy, risk, and governance for many years. Indeed we have been involved at a core level from the beginning, two years ago, in the work of the International Standards Organisation (ISO) on AI standards.

This AI standards work is undertaken within a body called 'Joint Technical Committee 1, Sub Committee 42 (JTC 1/SC 42)[5]. Microsoft's contribution to this effort in JTC 1/SC 42 is, we believe, unrivalled by any other company working on AI. At its most recent plenary in Japan in October 2019, nine Microsoft technical experts[6] attended with three in the role of editors/project leads for emerging new specifications. The reason for our work here is because we believe that standards provide an important means to give technical effect to a principle or regulation in a way that can be evaluated, compared, and most importantly certified.

[4] There are many consortiums that produce specifications, several of which have merit; a standard produced via the main standards-setting bodies, including ISO, is underpinned by a more rigorous process.

[5] https://www.iso.org/committee/6794475.html, last accessed 5 November 2019.

[6] Among the Ireland delegates in attendance was Microsoft's Terry Landers.

Increasingly, the standards work is not just technical, but also must reflect the evolving policy, legislative, and regulatory landscape for AI. For example, recent standards work has considered the EU General Data Protection Regulation (GDPR) and the Network and Information Systems Directive(NIS), the OECD and ITU initiatives on AI, and on the work of the EU Commission High Level Expert Group on AI.

Principles are incredibly helpful for designing legislative frameworks, but a standard builds on principles and provides further assurance to the consumer. By developing and adopting standards we provide a compliance pathway to best practices.

The international standards system can provide a useful space for dialogue between these stakeholders around the world, which reconcile the needs for compliance and innovation, and which we believe can lead to improved understanding of the key challenges and how they can be addressed. See Figure 1.

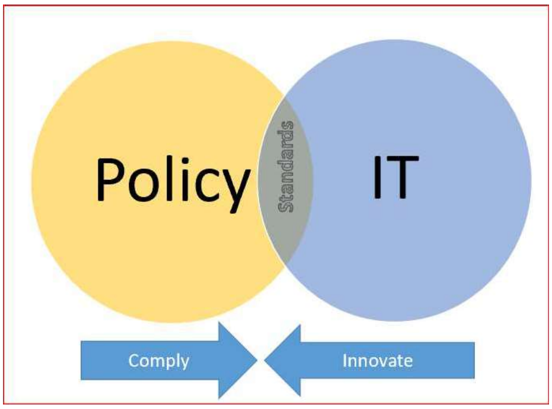

Figure 1

4.2. AI Standards development (Existing Standards Efforts)

To date, the standards efforts on AI have focused on the following topics:

4.2.1. Governance;
4.2.2. Bias;
4.2.3. Trustworthiness; and
4.2.4. Transparency.

4.3. AI Governance Standard (Under Development)

This draft standard, which builds on the existing IT and Data Governance standards[7], will provide guidelines and governance requirements that address the following issues:

4.3.1. Culture and values;
4.3.2. Compliance;
4.3.3. Risk framework;
4.3.4. Sources of risk; and
4.3.5. Controls.

4.4. Bias in AI Systems and AI-aided Decision-making (Under Development)

Bias can arise in many forms, including personal bias (unfounded intuition), outdated facts, skewed information, seeking causality, and illusion of confidence; these can lead to in-correct decision-making.

This draft standard[8] addresses areas of potential bias, including:

4.4.1. Cognitive bias (including automation bias and confirmation bias);
4.4.2. Statistical bias (for example, selection, sampling, and coverage);
4.4.3. Data bias (for example, gender, race, age); and
4.4.4. Bias in the model architecture.

It also addresses measures that may be used in order to mitigate bias in AI systems.

4.5. AI Trustworthiness[9]

In ISO, trust has been defined as '*the degree to which a user or other stakeholder has confidence that a product or system will behave as intended*'[10]. It reflects a sense of reliability, predictability, confidence, and compliance, but crucially reflected from the perspective of the user, service customer, or even regulator. The

[7] IS 38500 and IS 38505.

[8] ISO 24027:2019(XDraft) Artificial Intelligence (AI) – Bias in AI systems and AI aided decision making.

[9] ISO 24028:2019(X) Artificial intelligence (AI) — Overview of trustworthiness in Artificial intelligence.

[10] Source: ISO 25010:2011; Clause 4.1.3.2

challenge for service-providers is to engineer such trustworthy products and services.

Already, there is considerable work underway on trustworthy IT within ISO, but AI represents an additional unique set of trust scenarios in IT. The work on trustworthy AI recognises a range of potential AI vulnerabilities and potential measures to mitigate them, including:

4.5.1. vulnerabilities that can include: new security threats (since AI can be considered a 'dual -use technology'[11]), new privacy threats, and challenges related to the development and deployment of AI systems; and

4.5.2. mitigation measures to include 'human-in-the-loop' oversight, controllability, and explainability, for example, the use of counter-factuals[12].

4.6. AI and Transparency

Article 5 of the GDPR sets out a how data controllers must process all data: *"lawfully, fairly and in a transparent manner in relation to the data subject".* ISO's Compliance Management system standard, ISO 19600, provides guidance for establishing, developing, implementing, evaluating, maintaining, and improving an effective compliance management system; it is based on the principles of good governance, proportionality, transparency and sustainability.

[11] Can be used for both attack and defence.

[12] Example: someone who has been refused a bank loan would be advised the detailed reasons why and what are the main drivers to secure a positive response to a loan application in the future.

5. Learning From International Best Practice

5.1. Over the past 24 months we have seen many countries around the world define their '*National AI Strategy'* [13]; these documents can provide a useful reference for Ireland as it develops its own National AI strategy. As a reference, we have included a comparative summary in Appendix 1, from which we have deduced the following learnings that can help inform the strategy and the correlating actions to support the strategy:

5.1.1. Preparing Society

Preparing society for the socio-economic changes that will likely arise from AI disruption, especially in the labour market, when, for example driverless vehicles, or AI-based medical diagnostics enter the mainstream. This effort should seek to inform, to engage, and to empower people to embark on their own change journey.

5.1.2. Building Capacity

A pipeline of talented people skilled in AI will be a key requirement across all these domains. That requires us to understand the needs, and to address them in education, research (e.g., foundational AI research but also applied AI research), across the enterprise sector (e.g. existing enterprises and new AI-based enterprises) and in public service, for example, in healthcare, economic management, environment and law enforcement.

5.1.3. Supporting Infrastructure

Recognising the need to revise existing policy, regulation, and legislation, or to create new instruments for the emerging AI scenarios. It should also consider how we balance the benefits of AI, while mitigating the associated risks, and on the need to involve societal stakeholders, such as professional bodies (e.g., in law, accounting, sciences), trade unions, public service, and wider civic society in the conversation. It might also cover adjacent policy issues of privacy, data protection/open data, ethics, transparency, and governance as they relate to AI.

5.1.4. Creating Exemplars

Identifying and incubating Irish best-practice exemplars and use-cases that

[13] https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd ; http://www.oecd.org/going-digital/ai/initiatives-worldwide/, last accessed 6 November 2019.

can be show-cased, referenced and re-used by others.

5.1.5. Investment

For Ireland to fully realise the potential benefits of AI for our communities and our economy, some incremental strategic investments will be required. We see many examples around the world of how even smaller countries are making these investments, which should be commensurate with the scale of ambition we have for AI in Ireland.

5.1.6. Leadership from the Top

To drive execution of the national agenda, to manage resourcing, and work in collaboration with our partners in the European Union and elsewhere, in order to ensure that Europe's distinctive culture and values are reflected in AI we use every day.

Summary of Recommendations

1. Ireland should position itself as an early influencer of EU legislation for the regulation of facial recognition technologies, demonstrating leadership, expertise, and agility required to guide the responsible use of rapidly changing technology, but also as a strong advocate of fundamental human rights.

2. A new National Strategy on AI should seek to define principles in law consistent with European values and fundamental rights including:

   a. strengthening existing laws and promoting the development and deployment of AI systems that are lawful, ethical, and robust;
   b. creating incentives for the private sector to develop and deploy AI systems that reflect European values; and
   c. removing barriers to the European single market and encourage innovations that contribute to the social good.

3. The Government should ensure the new strategy both in Ireland and across the EU provides for appropriate legal and ethical protections against sensitive use cases, including:

   a. introducing safeguards against AI-based solutions that pose a material risk of consequential impact on individuals and society, such as facial recognition technology;
   b. incentivising the use of documentation models of data sets while at the same time guaranteeing flexible processes, formats, and tools;
   c. introducing frameworks for ensuring that facial recognition services are subject to third parties testing and publishing reasonable tests of these services for accuracy and unfair bias;
   d. strengthening privacy legislations by analysing the tension between privacy and the use of data for societally critical uses of AI; and
   e. limiting the ability to use facial recognition technologies for mass surveillance purposes.

4. Ireland already exerts influence in the international standardisation effort on AI. We should seek to share and leverage that know-how in order to educate key stakeholders about AI issues, and to drive the adoption of responsible AI underpinned by AI standards. We should also aspire to help shape the evolving standards agenda, especially on Trusted AI.

5. The Government should also ensure its own law and policymakers are building the capability to bridge technical and policy expertise in order to produce quality legislation as new technology emerges to harness the benefits and control the risks.

Closing

Once again, Microsoft thanks the Department of Business, Enterprise and Innovation for the opportunity to provide our input to this consultation and we welcome follow-up discussions at the Department's convenience.

Appendix 1
National AI Plans Compared

Canada	China	EU	Finland
Increase the number of AI researchers and graduates	(1) Focus on developing intelligent & networked products such as vehicles, service robots, & identification systems	(1) Increase the EU's technological and industrial capacity and AI uptake by the public and private sectors;	We will enhance Competitiveness of companies through the use of AI
Establish three clusters of scientific excellence	(2) Emphasize the development AI's support system, including intelligent sensors and neural network chips	(2) Prepare Europeans for the socioeconomic changes brought about by AI;	We will utilise data in all sectors
Develop thought leadership on the economic, ethical, policy, and legal implications of AI	(3) Encourage the development of intelligent manufacturing,	(3) ensure that an appropriate ethical and legal framework is in place.	We will speed up and simplify the adoption of artificial intelligence
Support the national research community on AI.	(4) Improve the environment for the development of AI by investing in industry training resources, standard testing, and cybersecurity		We will ensure top-level expertise and attract top experts

New Zealand	Singapore	Estonia	Rep. of Korea
(1) developing a coordinated national AI strategy,	Research: Invest in deep capabilities to catch the next wave of scientific innovations and breakthroughs	ADVANCING THE UPTAKE OF AI IN PUBLIC SECTOR IN ESTONIA • 30 supporting activities	First, to secure AI talent, the government will establish six graduate school in AI by 2022
(2) creating awareness and understanding of AI in the public,	Technology: Invest in deep capabilities to catch the next wave of scientific innovations and breakthroughs	ADVANCING THE UPTAKE OF AI IN PRIVATE SECTOR IN ESTONIA • 12 supporting activities	2. Development of AI technology. The government will fund large scale projects in national defence, medicine, and public safety
(3) assisting the public and private sectors adopt AI technologies,	Innovation: Broaden the use and adoption of AI in Singapore and groom local AI talents to support industry growth	DEVELOPING AI R&D AND EDUCATION IN ESTONIA • 9 supporting activities	3. the government will invest in infrastructure to support the development of AI start-ups and SMEs
(4) increasing access to trusted data,		DEVELOPING LEGAL ENVIRONMENT FOR UPTAKE OF AI 1 supporting activity	